Eco514—Game Theory
Lecture 15: Sequential Equilibrium

Marciano Siniscalchi

November 11, 1999

Introduction

The theory of extensive games is built upon a key notion, that of *sequential rationality*, and a key insight, the centrality of *off-equilibrium beliefs*. The definition of sequential equilibrium brings both to the fore in a straightforward manner, and emphasizes their interrelation.

From subgame perfection to sequential rationality

.

Let us begin by considering a game Γ with observed actions (but possibly simultaneous moves). Fix a profile $s \in S$; according to the definition, s is a subgame-perfect equilibrium of Γ iff, for every history $h \in H \setminus Z$, the continuation strategy profile $s|_h$ is a Nash equilibrium of the subgame $\Gamma(h)$ beginning at h.

Equivalently, we can say that s is a SPE iff, for every history $h \in H \setminus Z$, and every player $i \in N$, there is no strategy $s'_i \in S_i$ that yields a higher payoff to Player i than s_i in the subgame $\Gamma(h)$.

Moreover, we can actually restrict this requirement to subgames starting at histories $h \in P^{-1}(i)$. Consider a history h' such that $i \notin P(h')$; then either Player i never moves in the subgame $\Gamma(h')$ (in which case s_i is trivially optimal), or there exists a collection of histories $\{h''\}$ such that (i) h' is a subhistory of every h'', and (ii) $i \in P(h'')$ for each such h''. According to the profile s_{-i}, exactly one of these histories, say h, will be reached starting from h', and by assumption $s_i|_h$ is a best reply to $s_{-i}|_h$ in $\Gamma(h)$. Therefore, $s_i|_{h'}$ is also a best reply to $s_{-i}|_{h'}$ in $\Gamma(h')$.

In other words, s is a SPE iff each component strategy s_i is a best reply against s_{-i} *conditional upon reaching any history where Player i moves*—including those histories that are not generated by s, or perhaps explicitly excluded by s_i itself. This is precisely the notion of *sequential rationality*.

VISIT…

There is little reason why this concept should not apply to general extensive games as well. However, observe that, in games with observed actions, fixing a strategy profile s is sufficient to generate well-defined beliefs in every subgame—i.e. well-defined beliefs conditional upon reaching any history. This is not the case in general extensive games (or in Bayesian extensive games with observed actions, for that matter).

In some (very special) games, this does not really matter. For instance, consider Fig. 1.

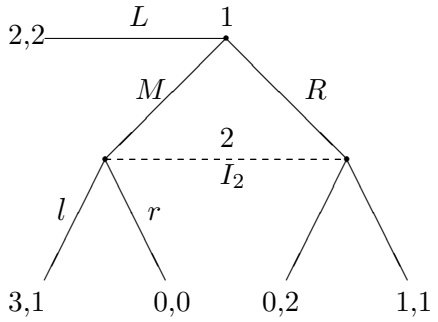

Figure 1: A special game (OR 219.1)

Observe that (L, r) is a Nash equilibrium of this game; note also that, since Player 1 chooses L in this equilibrium, it is not clear what Player 2 should believe if the information set $I_2 = \{(M), (R)\}$ is reached! Thus, *the equilibrium strategy profile is insufficient to pin down players' conditional beliefs.*

Yet, in this particular game, this is essentially irrelevant: regardless of his beliefs about Player 1's actual choice, r is conditionally strictly dominated at I_2. Whatever Player 2's beliefs, r cannot thus be sequentially rational, so (L, r) is "not a legitimate solution" to the extensive game in Figure 1. On the other hand, (M, l) is.

In this simple game, Player 2's beliefs about the strategy (action) actually chosen by Player 1 is irrelevant as far as his choices are concerned. However, this is seldom the case: consider Fig. 2, which differs from Fig. 1 only in that the payoffs at the terminal history $\{M, r\}$ have been modified.

Consider (L, r) again. Upon being reached, Player 2 must conclude that Player 1 did not choose L; hence, he must adopt new, *distinct from the equilibrium profile.* Suppose that he thinks that Player 1 actually chose M: then r is indeed a sequential best reply! Suppose instead that he thinks that Player 1 chose R: in this case r is not sequentially rational.

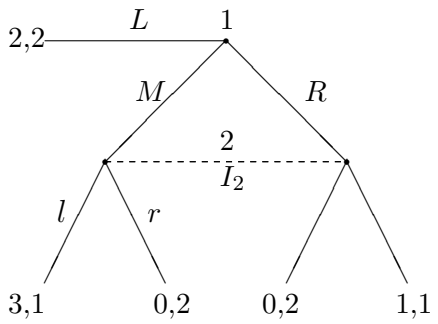

Figure 2: A typical case (OR 220.1)

Formal definitions

The simple games in the previous section emphasize that *off-equilibrium beliefs matter*: they determine whether or not a given profile is an equilibrium. Moreover, *the putative equilibrium profile does not yield a complete specification of beliefs.*

In particular, at any information set, a (behavioral) strategy profile β yields a complete specification of continuation strategies, but may fail to provide information about past play—precisely because, at information sets off the path of play generated by β, it must necessarily be the case that at least one player j made choices other than those prescribed by β_j! The original profile is clearly silent about such alternative choices.

Note that this problem does not arise in games with observed actions, precisely because past play is observable (so there is no uncertainty about previous moves). However, it *does* arise in Bayesian games with observed actions; see the next section for details.

Kreps and Wilson propose the following definition.

Definition 1 *Fix a general extensive game Γ. A belief system is a map $\mu : \bigcup_{i \in N} \mathcal{I}_i \to \Delta(H \setminus Z)$ such that, for every $i \in N$ and $I_i \in \mathcal{I}_i$, $\mu(I_i)(I_i) = 1$.*

It is most helpful to view a belief system as a *summary representation of past play leading to information sets.* That is, informally speaking, one could imagine that, if the play reaches an information set I_i (where Player i has to move) which is not consistent with the originally expected behavioral strategy profile β, players conclude that a different profile β' is actually being played, and base their inferences and forecasts on the latter.

However, based on simple consistency considerations, β' and β should agree at least as far as moves at information sets following I are concerned. But, if this is the case, then all that is required in addition to β in order to verify the sequential rationality of β_i at I_i is the distribution over histories in I_i induced by β'—which is precisely the type of information encoded in $\mu(I)$.

From now on, the basic unit of our analysis will be a pair (β, μ) consisting of a behavioral strategy profile and a belief system; such a pair is called an *assessment*.

Focus on games with perfect recall. In keeping with conventional usage, we define sequentially rational assessments. First, define the *conditional outcome function* $O(\beta, \mu|I)$ to be the distribution over Z induced by the assessment (β, μ) conditional upon starting the game at I: formally, for every $z \in Z$, if no $h \in I$ is a subhistory of z, then $O(\beta, \mu|I)(z) = 0$; otherwise, by perfect recall, there can be only one such subhistory h (if there were two, then one would have to be a subhistory of the other, and an information set cannot contain a pair of nested histories). Thus, $z = (h, a^1, \ldots, a^K)$, and we let

$$O(\beta, \mu|I)(z) = \mu(I)(h) \prod_{\ell=1}^{K-1} \beta_{P(h,a^1,\ldots,a^\ell)}(h, a^1, \ldots, a^\ell)(a^{\ell+1})$$

(recall that, with some abuse of notation, for $h \in I_i \in \mathcal{I}_i$, $\beta_i(h) = \beta_i(I_i)$).

Definition 2 *Fix a general extensive game Γ. An assessment (β, μ) for Γ is sequentially rational iff, for every player $i \in N$, every information set $I_i \in \mathcal{I}_i$, and every behavioral strategy $\beta_i' \in \mathcal{B}_i$,*

$$\sum_{z \in Z} U_i(z) O(\beta, \mu|I_i) \geq \sum_{z \in Z} U_i(z) O((\beta_i', \beta_{-i})\mu|I_i)$$

It should be obvious that *some* relationship between belief systems and behavioral strategies in an assessment (β, μ) should be maintained.

In particular, at least at information sets consistent with β, μ should be derived from β by means of Bayes' rule.

But we may reasonably require more than this. For instance, consider a modification of the game in Figure 2 in which Player 1 chooses between L and C, and if 1 chooses C, a third player $1'$ chooses between M and R; then Player 2 has to choose observing only Player 1's initial move. That is, Player 1's choice among L, M and R is split between two players, 1 and $1'$, choosing consecutively.

Consider a behavioral strategy profile in which Player 1 chooses L and Player $1'$ chooses M. Then, *even after a deviation by Player 1*, we would probably still want to require that $\mu(\{(C, M), (C, R)\})((C, M)) = 1$; that is, a deviation by Player 1 should not affect Player 2's beliefs about the choices of Player $1'$. However, strictly speaking, Bayes' rule does not apply to this information set, because it lies off the path induced by the behavioral strategy profile under consideration.

[Note incidentally that if 1 and $1'$ were the same player, this argument would be less convincing; one would have to invoke some kind of *trembles* or mistakes, whereas the argument in the case where 1 and $1'$ are distinct is essentially based on *independence*. As will be clear, however, the standard definition of sequential equilibrium does *not* discriminate between these two cases.]

Kreps and Wilson suggest the following definition.

Definition 3 *Fix a finite general extensive game Γ. An assessment (β, μ) is consistent iff there exists a sequence of assessments (β^n, μ^n) converging to (β, μ) and such that (i) $\beta_i^n(I_i)(a) > 0$ for every $i \in N$, $I \in \mathcal{I}_i$ and $a \in A(I)$; and (ii) μ^n is derived from β^n via Bayes' rule.*

It is easy to see that this definition captures the intended restrictions. Note that (ii) makes sense because, since every action in the game receives positive probability, every information set is also reached with positive probability, so Bayes' rule always applies.

Myerson suggests the following argument. Almost all elements of the set $\mathcal{B}$ of behavioral strategy profiles satisfy Condition (i) in Definition 3: that is, Condition (i) is generically satisfied. Also, Condition (ii) is very reasonable, so we definitely want all assessments (β, μ) such that β satisfies (i) to be deemed consistent if and only if μ is derived from β by Bayes' rule. Finally, it seems reasonable (and convenient) to assume that the set of consistent assessments is closed; hence, it makes sense to define it to be precisely the closure of the set of assessments satisfying (i) and (ii).

This may not appear to be a tremendously compelling argument; however, let me add that, for certain classes of games, consistency is equivalent to much simpler conditions on beliefs (when the latter are represented in an appropriate format). Also, consistency can be characterized for general games in a manner that does not involve limits (although the characterization is not tremendously appealing).

We are finally ready to define sequential equilibrium.

Definition 4 *Fix a finite general extensive game Γ. An assessment (β, μ) is a sequential equilibrium of Γ iff it is consistent and sequentially rational.*

Perfect Bayesian equilibrium

Some people find the definition of sequential equilibrium unwieldy (this refers in particular to the consistency requirement). Also, consistency is only defined for finite games, so sequential equilibrium does not really apply to relatively simple games such as Bayesian extensive games with observed actions where type sets are infinite.

In such games, just like in a Bayesian game with simultaneous moves, one must specify a strategy for each payoff-type of each player. That is, we must specify type-contingent strategies. Hence, upon reaching a history h, players may be able to make inferences about their opponents' types based on the (equilibrium) strategy profile. However, at histories off the anticipated path of play, it is clear that the strategy profile does not convey any information about opponents' types.

This is of course exactly the type of problem that beliefs system solve in general games. Here, however, the situation is much simpler, because the domain of uncertainty (the sets of payoff types) are fixed throughout the game.

In Perfect Bayesian equilibria (see OR for details), the information conveyed by strategy profiles is complemented by measures $\mu_i(h) \in \Delta(\Theta_i)$ associated with each nonterminal history.

In interpreting the definitions, it is important to remember that, just like the probabilities p_i, $\mu_i(h)$ represents *the beliefs of players other than i about Player i's type*. These beliefs are common to all players, which is reasonable: after all, the priors p_i are common, and all players observe exactly the same occurrences as the play progresses.

The key condition in the definition of a PBE is the requirement that these beliefs be *action-determined*. That is: if Player i does not move at a history h, then the beliefs about her should not change at h; and if Player i does move, than beliefs about her should only be affected by her own action at h (not by somebody else's action).

Just like the consistency requirement in SE, this complements Bayesian updating with restrictions on beliefs following unexpected actions (i.e. on beliefs at off-equilibrium histories).